

# **NEXT STEPS IN NETWORK TIME SYNCHRONIZATION FOR NAVY SHIPBOARD APPLICATIONS**

**Karen O'Donoghue, Mark Glass, and Timothy Plunkett**  
**NAVSEA Surface Warfare Center, Dahlgren Division**  
**17214 Avenue B, Suite 126**  
**Dahlgren, VA 22448-5147, USA**  
**Tel: +1-540-653-1567**

## **Abstract**

*Next generation Navy platform designs are evolving towards generalized multipurpose infrastructures based on open standards and commercial products. These platforms will support a wide range of new and expanding applications in a more flexible and dynamic manner than in previous designs. This new paradigm creates significant network time synchronization challenges. The Navy has been deploying the Network Time Protocol (NTP) in shipboard computing infrastructures to meet the current network time synchronization requirements. Additionally, a new standard, IEEE 1588 or the Precision Time Protocol (PTP), has emerged. It holds the promise of more precise synchronization through the use of hardware assists. New international standardization efforts intend to leverage NTP and PTP for a next generation of time synchronization protocols. This paper examines the emerging work in the International Standards communities in the area of precise network time synchronization. Additionally, it looks at preliminary evaluations of NTP and PTP in the types of products commonly used in Navy shipboard applications. Finally, this paper proposes potential focus areas for Navy standardization and experimentation efforts.*

## **INTRODUCTION**

Navy shipboard applications are inherently time critical. The calculations involved in the detection, location, and identification of a moving target made with respect to a moving platform requires precise timing. These requirements were historically met using Navy specialized computers and networks. The Navy has evolved to an Open Architecture approach using commercial networks, operating systems, and processors. As a result, standards-based approaches for network time synchronization are being used. As applications and the underlying computing environments become more complex and distributed, the current requirements and solutions will be insufficient to meet the next generation time critical application requirements.

## **KEY REQUIREMENT DRIVERS**

Several key drivers have been identified that will impact the requirements of the next generation time critical applications. The first driver is the evolving application domain. There is a basic sequence of events that occur in a Navy platform. First, the target must be detected. Then, it must be identified and tracked. Finally, the target may be

## Report Documentation Page

*Form Approved*  
*OMB No. 0704-0188*

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE <b>01 DEC 2008</b>	2. REPORT TYPE <b>N/A</b>	3. DATES COVERED <b>-</b>	
4. TITLE AND SUBTITLE <b>Next Steps In Network Time Synchronization For Navy Shipboard Applications</b>		5a. CONTRACT NUMBER	
		5b. GRANT NUMBER	
		5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)		5d. PROJECT NUMBER	
		5e. TASK NUMBER	
		5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) <b>NAVSEA Surface Warfare Center, Dahlgren Division Dahlgren, VA 22448-5147, USA</b>		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSOR/MONITOR'S ACRONYM(S)	
		11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT <b>Approved for public release, distribution unlimited</b>			
13. SUPPLEMENTARY NOTES <b>See also ADM002186. Annual Precise Time and Time Interval Systems and Applications Meeting (40th) Held in Reston, Virginia on 1-4 December 2008, The original document contains color images.</b>			
14. ABSTRACT			
15. SUBJECT TERMS			
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT
a. REPORT <b>unclassified</b>	b. ABSTRACT <b>unclassified</b>	c. THIS PAGE <b>unclassified</b>	<b>UU</b>
			18. NUMBER OF PAGES <b>10</b>
			19a. NAME OF RESPONSIBLE PERSON

engaged or monitored over a period of time. Historically, this sequence of events has taken place on a single platform within a single or a small set of sensors and processors with the weapon or monitoring system used for engagement located on that same platform. In the future, these sensors and processors may be located on multiple independent platforms with the weapon or monitoring system used to engage the target on yet another independent platform. This distribution of the participating components has obvious consequences for the role that the network time synchronization will play in the successful engagement of targets in next generation shipboard combat systems.

The second driver is the evolution of the computing infrastructures used to deploy these applications. In the past, these infrastructures were tightly coupled to the applications that ran on them. The future direction of combat systems proposes these infrastructures be partitioned from the applications and be designed and procured independently. Figure 1 provides a generalization of this concept. Implementation of this concept requires the computing infrastructure to be open and standards-based. In addition, mainstream solutions will take precedence over niche solutions. History has shown that mainstream commercial products are more likely to evolve gracefully into the next generation products available to Navy Program Managers.




Figure 1. Conceptualized Navy open system implementation.

The third driver is the scale of the systems involved. It is expected that the number of sensors and processors participating in a particular engagement scenario will increase dramatically with the distribution of the application. Current systems contain tens of components, whereas future systems may have hundreds of participating components.

Another driver is the increasing importance of high-fidelity data recording and tightly integrated virtual training systems. In an era of cost and manpower conservation, the relevance and required fidelity of these systems is increasing.

In summary, the targets will be more challenging, there will be more components involved in target engagement, and these components will be more widely dispersed. All these drivers have significant implications for the next generation network time synchronization requirements.

## EVOLVING COMPUTING INFRASTRUCTURES

Navy shipboard systems have evolved from a relatively small number of Navy specified computers interconnected by specialized high bandwidth Navy I/O channels to a large number of commercial processors interconnected by standards based commercial networks. Time services were originally provided by specialized time processor hardware and a separate hardware-based time distribution system that provided time pulses to each tactical computer. As Navy shipboard systems have evolved towards commercial processors and standards-based networks, Navy time services have moved away from a separate time distribution system and towards utilization of the same networks used for communications. This movement creates a challenge to meet stringent time synchronization performance requirements while utilizing shared commercial processor and network resources.

Figure 2 illustrates the Network Time Synchronization Architecture for a notional Navy shipboard combat system under development today. Time and position information is received from Global Positioning System (GPS) satellites by a shipboard Common Time Reference (CTR) system. An Inter-Range Instrumentation Group (IRIG) signal is sent to the Time Processor from the CTR. The Time Processor reads the IRIG time and sets the local operating system time based on the information received from the CTR. A Network Time Protocol (NTP) server is running on each Time Processor. All tactical nodes in the system are NTP clients of the two Time Processors. In this scenario, a separate time manager monitors the relationship between the Forward and AFT time processors such that the two stay within a specified tolerance of each other. The Network Time Synchronization Architecture represents a major step in the migration from the use of specialized Navy computers and separate time distribution networks.

The next step beyond the architecture shown in Figure 2 is to eliminate the IRIG time distribution network and specialized time processors. Figure 3 shows a notional next generation architecture where the CTR receives the reference time from GPS and provides time to all the tactical processors over the combat system LAN using a commercial implementation of a standard network time synchronization protocol. In this scenario, non-tactical and tactical systems will all be getting time information from the same set of coordinated Common Time References.

As the Navy has evolved towards the use of commercial products based on open standards, it is clear that the Navy needs to understand and influence necessary both the developing standards and the emerging products. The main objective during this effort is to ensure that cost-effective products that meet Navy requirements are available in the marketplace. A number of hard lessons have been learned along the way regarding the use of open standards and the availability of products which will help the Navy in meeting this objective. First, widely used commercial standards, as opposed to Navy unique standards, must be utilized. Second, the proprietary features or obscure options must be avoided. A capability in a standards document does not guarantee a feature generally available in a product. Two general types of activities will help the Navy further facilitate this objective. First, the Navy must ensure that emerging open standards for network time synchronization are sufficient to meet next generation requirements. Additionally, the Navy must examine emerging products for performance and capability gaps. The following two sections discuss current and proposed efforts in the areas of standardization and experimentation.




Figure 2. Notional Current Navy network time synchronization architecture.




Figure 3. Notional next-generation Navy network time synchronization architecture.

## STANDARDIZATION ACTIVITIES

Network time synchronization protocols utilize an existing communications infrastructure to exchange time information for the purposes of time synchronization. The current most widely deployed network time synchronization protocol is the NTP. NTP has evolved over the last 20 years as the primary technology for time synchronization over Internet Protocol (IP)-based networks, in particular for wide area networks such as the Internet [1]. The Navy has invested a significant amount of development and analysis into the deployment of NTP in Navy systems. NTP is now the predominant mechanism for network time synchronization in Navy shipboard systems. The current published specification, which documents NTP Version 3, is RFC 1305 [2]. A subset of NTP for simple clients referred to as Simple NTP (SNTP) is specified in RFC 4330 [3]. There is an ongoing effort in the Internet Engineering Task Force (IETF) to complete the NTP Version 4 (NTPv4) specification. This work has been completed by the IETF NTP working group and is in the final stages of approval and publication. Publication is expected in the middle of 2009. This update, which incorporates both RFC 1305 and RFC 4330, includes numerous improvements including faster initial convergence times, improved algorithms, and support for the Internet Protocol Version 6 (IPv6).

Despite the success of NTP, there are concerns about the next generation time requirements and the ability of NTP to evolve to meet these requirements. Because of this concern, the Navy has been following the development of a new time synchronization protocol, IEEE 1588, also known as the Precise Time Protocol (PTP). IEEE 1588 Version 1 was developed primarily by the test and measurement industry with participation from the industrial automation community [4] and was published in 2002 [5]. IEEE 1588 Version 2 [6] is a significant update which was developed with broader industry participation. PTP is similar to NTP in many ways; however, it was developed with a focus on precise time synchronization in more focused geographically environments similar to those found on Navy shipboard platforms. Additionally, PTP was designed with the expectation that basic time stamping capability would be built into commercial network hardware. Thus, it holds the promise of time synchronization with hardware support for future Navy systems.

With the evolution of PTP, the Navy is looking to define a coordinated time architecture that leverages its investment in NTP while positioning it to utilize next generation hardware and software techniques to meet emerging requirements. It is expected that the Navy's next generation shipboard time synchronization service will be an evolution of the current NTP and PTP solutions. There are a number of key ongoing international standardization activities related to network time synchronization that are of significant interest to the Navy.

The first of these activities is the IETF. The IETF is the international standards organization responsible for all of the standards used in the global Internet. The IETF is finalizing the work of its NTP WG and has, as of March 2008, chartered a new effort to investigate next generation network time synchronization requirements. This new effort, the Transmitting Timing over IP Connections and Transfer of Clock (TICTOC) working group, is currently collecting timing requirements from various application areas. When the requirements effort has concluded, the TICTOC WG will analyze the current solutions in the context of these requirements and propose a way ahead. It is expected that this way ahead will address both IEEE 1588 and NTP. There are a number of additional work items to make IEEE 1588 more robust in a broader wide area network context. There are also a number of possible enhancements to NTP, including faster polling intervals, standardized follow-up message capabilities, and the use of alternative clock algorithms. Further information on the IETF and the efforts of the TICTOC working group can be found at [www.ietf.org](http://www.ietf.org).

The second activity of interest is the IEEE 802.1 standards activity sponsored by the IEEE Standards Association. The IEEE 802 LAN/MAN Standards Committee is the international standards organization responsible for the development of Local Area Network (LAN) and Metropolitan Area Network (MAN) standards. This includes all the commonly used wired and wireless network technologies including Ethernet, Wi-Fi, and WiMax. The IEEE 802 time synchronization effort evolved out of an activity to provide better support for distributed audio/video applications in the home. The resulting standards development effort, *IEEE 802.1AS Standard for Time Sensitive Applications in*

*Bridged LANs*, is being developed by the IEEE 802.1 Task Group. The initial Project Authorization Request (PAR) was approved in July 2006. An early decision was made in this effort to base its time synchronization mechanisms on the IEEE 1588 standard. This had obvious benefits to the Navy in the sense that it broadened the source of potential IEEE 1588 products while reducing the number of similar possible conflicting solutions. The possibility that, at some point in the future, all IEEE 802 products could support time synchronization is very appealing. Further information on IEEE 802 activities can be found at [www.ieee802.org](http://www.ieee802.org).

A third activity of interest is the International Telecommunication Union (ITU) Telecommunication Standardization Sector (ITU-T). ITU-T has many years of experience specifying synchronization for Time Division Multiplex (TDM) networks. As telecommunication networks have migrated away from TDM networks towards packet-based networks, ITU-T has started to study the transport of synchronization through these packet networks. ITU-T SG 15 Q13 (Question 13 of Study Group 15 of ITU-T) has completed definitions of jitter and wander for circuit emulation services and the specification of Synchronous Ethernet. It is currently working on clock specifications for packet networks and metrics for characterization of those networks. This work includes the definition of a telecoms profile for IEEE 1588v2. There is an active effort to ensure that the IEEE, IETF, and ITU-T activities are closely coordinated. Further information on ITU-T activities can be found at [www.itu.int/ITU-T/](http://www.itu.int/ITU-T/).

## **EXPERIMENTATION ACTIVITIES**

As discussed above, the Navy has extensive experience with NTP and is evaluating the use of PTP for deployment. A series of experiments have been proposed to investigate the capabilities and limitations of emerging products. The primary objective of this analysis is to gather data on the performance of hardware- and software-based PTP solutions. As a first step in this analysis, two basic experimental questions were proposed. First, what is the relative difference in the performance of NTP and PTP when deployed in the same notional Navy shipboard system? Second, what performance improvements are demonstrated in a notional Navy system when hardware assists are available? The resulting experiments are described in some detail in [7] and summarized below.

### **EVALUATION OF NTP AND PTP IN SOFTWARE**

In this experiment, software implementations of NTP and PTP were compared in an experimental configuration similar to the system architecture to be used onboard Navy platforms and illustrated earlier in Figure 2. In this experiment, the network switches were standard switches that were not PTP aware, which means that they simply passed PTP packets like any other network traffic. All the computer nodes in this experiment were single-board computers with multi-core processors running a real-time Linux operating system.

One computer node served as an NTP server. It received time via an IRIG input and distributed the time across the computer network via the NTP protocol. A second computer node was configured to be an NTP client. This node used NTP to synchronize its local computer clock. This computer node was also outfitted with an IRIG card so that the local computer time could be compared with the time received via IRIG.

In order to gather PTP time in an identical configuration, another computer node was added to serve as the PTP master. The NTP client was converted into a PTP slave by stopping the NTP process and running the PTP software implementation as a PTP slave. Initial results showed very poor synchronization and, as a result, the PTP process was isolated to its own CPU core and elevated to run at highest system priority. This same procedure had no impact on the NTP client. The PTP software then synchronized the time on the local computer clock to the time received via the PTP protocol. Time offsets were measured by comparing the local computer time to the time received via IRIG. A summary of the results for both NTP and PTP is shown in Figure 4.

For this series of experiments, a software-only implementation of PTP performed better than a software-only implementation of NTP in a Navy notional shipboard computing environment with ideal conditions. Additional experiments are necessary to fully characterize PTP performance in this environment.




Figure 4. Software implementations of NTP and PTP.

## EVALUATION OF PRELIMINARY PTP HARDWARE IMPLEMENTATIONS

The second series of experiments focused on the performance improvements demonstrated in a notional Navy system when hardware assists are available. The time source is a GPS receiver that is capable of functioning as both an NTP server and a PTP master clock. NTP client and PTP slave functionality were both provided on the same non-real-time Linux PC. Two different network interfaces were used in the Linux PC, including a standard network interface card (NIC) and one with hardware to support PTP. All test data were obtained by comparing the time on the time source with the time available on the Linux PC through the use of a software measurement tool based on NTP. While this tool is not optimal for this analysis, it enabled an initial evaluation. The results of this evaluation can be used to define further experiments.

There were three basic variations executed for this test. In the first, NTP performance was measured over a network that made use of a standard network switch. The second variation involved synchronization of the clock on the PTP NIC with the PTP master clock using the PTP protocol over a network. Three network options were used in this variation, including a direct connection, a standard network switch, and a network switch that functioned as a PTP boundary clock. The last variation involved the use of PTP software that synchronized the Linux PC system clock by using the PTP protocol to obtain time information from the PTP master clock. The interconnecting network made use of both a standard network switch and the PTP boundary clock. The results for all these experiments are summarized in Figure 5.

Measurements for this experiment were made using a software tool for measuring clock offsets across the network. NTP offset measurements were made by measuring the offset between the local computer clock and the time provided by the NTP server. PTP offset measurements were made by comparing the time provided by the PTP card with the time provided by the NTP server (located on the same time server as the PTP grandmaster clock).




Figure 5. Hardware and software implementations of PTP.

In summary, for the above series of experiments, hardware assistance in both the NIC and the switch improved the performance of the PTP slave. In all cases, the PTP client performed better than the NTP client in this environment. When comparing the results of the first series of experiments with the second, it must be noted that different environments were utilized. In particular, the first environment utilized multi-core processors and a real-time operating system. Neither of these capabilities was utilized in the second series of experiments.

## CONCLUSION

The Navy is rapidly moving towards open systems, including commercial networks, to meet the computing environment needs of its real-time combat systems. A key aspect of this migration is the provision of a robust network time synchronization solution. Additional open standards development and experimental analysis are vital to fully realizing the Navy's future network time synchronization objectives.

Open standards activities key to next generation network time synchronization include the efforts of the IETF TICTOC WG, IEEE 802, and the ITU. These international standards activities are developing coordinated standards that will provide time synchronization in the computing products generally used by Navy systems. They are addressing a number of remaining open issues, including scaling, security, management, and metrics. The Navy needs to monitor and influence these activities where necessary to ensure that the Navy's emerging requirements can be met by the resulting standards.

Further experimental analysis is also required. First, a reevaluation of the tools and metrics used to analyze network time synchronization needs to be performed. The above series of experiments exposed the limitations of these tools, along with the necessity to develop more robust metrics for analysis. The emerging products may be more capable than the tools used to measure their performance. Second, a more robust analysis of PTP needs to be performed. This includes measurement of the performance of PTP in the presence of faults and heavy network loads, an evaluation of PTP on the various standard and real-time operating systems of interest to the Navy, and an analysis of performance of PTP as systems increase in scale. Execution of these experiments will help to insure that emerging products will be deployable in Navy systems.

The actions recommended in this paper will help to ensure that network time synchronization standards and products are available for Navy systems. Beyond that, the Navy still needs to develop candidate architectures for Navy use, provide best practices guidance to Navy Program Managers, and ensure that new time synchronization capabilities are reflected in standardized Application Programming Interfaces (APIs).

## ACKNOWLEDGMENTS

The authors wish to acknowledge Filoteo Garcia of Lockheed Martin, Jeff McDonald of Jtime, and Heiko Gerstung of Meinberg for providing loaner equipment that contributed to the execution of the experiments discussed above.

## REFERENCES

- [1] D. L. Mills, 2006, **Computer Network Time Synchronization** (CRC Press, Boca Raton, Florida).
- [2] D. Mills, 1992, RFC 1305 “*Network Time Protocol (Version 3) Specification, Implementation and Analysis*” (Internet Engineering Task Force).
- [3] D. Mills, 2006, RFC 4330 “*Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI*” (Internet Engineering Task Force).
- [4] J. C. Eidson, 2006, **Measurement, Control and Communication Using IEEE 1588** (Springer-Verlag, London).
- [5] IEEE Standard 1588-2002 “*IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*” (Institute of Electrical and Electronics Engineers, New York).
- [6] IEEE Standard 1588-2008 “*IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*” (Institute of Electrical and Electronics Engineers, New York).
- [7] M. E. Glass and K. F. O’Donoghue, 2008, “*Navy Shipboard Time Synchronization Service Options and Analysis*,” in Proceedings of the 2008 IEEE International Symposium on Precision Clock Synchronization (ISPCS), 22-26 September 2008, Ann Arbor, Michigan, USA (IEEE Publication), pp. 105-109.

